Disclaimer

Chapter 6, “Beam Diagnostics,” is not anywhere near completion. It is being released now because it contains some basic information about multiwires, BPMs, and BLMs. There are no diagrams. System experts have not recently evaluated it.

Future versions of this chapter will include information on flying wires and “remote sensing” instrumentation such as DCCTs, stripline detectors, and bunch length monitors. There will also be material on the timing of diagnostics, and, of course, pictures.

Chapter 6: Beam Diagnostics

> I’m beginning to wonder whether the physicists are pulling some kind of elaborate scam here. I’m starting to wonder if they don’t sit around their $23-million atomic accelerators all day, drinking frozen daiquiris and shrieking, “There goes one now!” and then laughing themselves sick. Maybe it’s time we laypersons asked some hard questions about this idea that all matter consists of tiny invisible particles whizzing around.
>
> - Dave Barry

Dave Barry is wrong of course, our accelerators cost a lot more than $23 million. But how do we know that there are tiny invisible particles whizzing around the Main Injector? And almost as importantly, how do we know that they are doing what they are supposed to? How wants a daiquiri? This chapter will attempt to answer those questions.

There are a few basic kinds of techniques used to look at beam. One requires that an object, such as a wire, be put directly into the beam. This method needs to be implemented with great care, because if done improperly the beam can be destroyed or significantly degraded. Devices called multiwires are used in "one-pass" beam lines, where the beam will only strike each set of multiwires once. Multiwires are inappropriate for looking at circulating beam, since the wires and the beam would be destroyed in very short order.

The solution adopted for using wires in a circulating beam is called a flying wire. The flying wires are very thin and move through the beam very quickly, so that disruption to the beam is minimized.

Both of these wire methods are destructive to the beam to some extent, so there is always a trade-off: Use the wires and degrade the beam, or don't use the wires and remain ignorant about what the beam is doing.

The second type of approach is less destructive. Protons and antiprotons have an electric charge, and can be "remotely" sensed by devices designed to detect beam electromagnetically. (This method is not entirely risk-free. There are image charges created in the instrumentation that can, in turn, influence the beam. There is no such thing as a free lunch.) Devices that detect the beam electromagnetically include the beam position monitors (BPMs), fast bunch integrators (FBIs), bunch length monitors, DC Current Transformers (DCCTs), SBDs (Sampled Bunch Displays), and resistive wall monitors.

A third, and indirect, way of "seeing" the beam is through the beam loss monitors (BLMs). When beam strikes an object, a spray of secondary particles is produced which can be detected outside of the beam pipe. (This is the reason no one is allowed in the tunnel when beam is present.) It is preferable that the losses not occur in the first place, but when they do, they provide valuable information about the location of the beam.

All of the methods mentioned above could be used to create a picture of what the beam is doing. Now, one at a time.

Multiwires

The multiwires are used in the beam lines: the MI-8 line, the abort line, and the P1, P2, P3, and A1 lines. This section describes the general principles behind the design and operation of multiwires; the specific layout in each line will be covered in Chapter 7.

When beam strikes a metal wire, electrons are dislodged. If the wire is part of a circuit, current will flow into the wire from elsewhere in the circuit to replace the missing electrons, and the charge imbalance can be measured. The greater the intensity of the beam striking a particular wire, the greater the charge displaced:

[Future Picture]

Horizontal and vertical sets of wires are strung in a G-10 board (G-10 is the name of a green fiberglass-reinforced epoxy commonly used at Fermilab). The wires are normally placed 1 mm apart. Sometimes the board is referred to as a paddle, because it rotates around the vertical shaft. The assembly is suspended in an evacuated metal can in the beam pipe. The horizontal set of wires measures the vertical profile of the beam, and vice-versa.

Because of the way that the board is rotated into the beam, only one plane can be measured at a time. The wires can move from the “out” position to either the horizontal or vertical position, and back again. Since the wires scatter some of the beam, it is usually preferable that the wires be out of the beam unless the beam profiles are actually being measured. Also, if the wires are in an intermediate position, the beam may hit the G10 board and not get through at all.

In addition to the position readbacks available in the control room, the position of the wires can be determined in the tunnel by looking at the top of the can. There is a pin visible from outside the can that is linked to the assembly:

[Future Picture]

The limit switches to either side of the “out” position provide the readbacks to the software read in the control room.

The cables that connect the multiwires to the electronics upstairs have a known capacitance, and cables that can store a charge:

[Future Picture]

When the time comes to read the beam intensity, the switches are closed and the charge accumulated by the wires is read by the electronics.

The result is a picture that looks something like this:

[Future Picture]

Each column in the histogram represents the beam intensity as measured by one wire. The horizontal and vertical planes are read separately.

The multiwires provide information about the shape of the beam that is not available from a simple position readback. For example, if the beam is too wide or too narrow, there could be a problem with one of the quadrupoles upstream; or, if the beam is scraping some object, the “shadow” can often be seen in the Multiwire plot. Once nominal profiles have been established, hardcopies or save files can be used to compare to current conditions and to diagnose problems.

Multiwire Controls

Some of the multiwires interface to the controls system through CAMAC. There are two types of CAMAC cards used by the multiwires. One is the 184 card, which is the motor controller card. It talks to the stepping motor power supply, which in turn sends streams of pulses to the motor downstairs that rotates the paddle. Each pulse rotates the motor through an incremental angle.

The other CAMAC card is the 192 card, which digitizes the data from the cables and makes it available to the CAMAC links. The 192 card has a small microprocessor on board which can make a few simple calculations on the data, including determination of the beam centroid (average center position) and the beam sigma (related to the width of the beam).

Multiwires not connected to CAMAC are operated through SWIC controllers similar to those used in Switchyard. They are linked through an ARCNET loop that feeds the data into the ARCNET card of a VME crate; the data returns to the computer room via Ethernet.

It is important that the wires sample the data at the exact moment that beam is present. The Multiwire electronics listens to the beam sync clocks to synchronize the sample time with the arrival of beam. (Remember that the beam sync clocks also trigger the kickers responsible for transferring beam.) Specific clock events used in each line are discussed in the chapter on beam lines.

Beam Position Monitors

The Beam Position Monitors, or BPMs, measure the average position of the beam at numerous locations in the beam lines and around the Main Injector ring. Unlike multiwires, which must be put directly into the beam, BPMs detect the beam electromagnetically. This property allows them to be used for measuring circulating beam.

The Main Injector BPMs are located just downstream (proton direction) of each quadrupole in the ring, and at strategic locations in each of the beam lines. The physical construction of the detectors differs from those of the former Main Ring in that each one can measure the beam position in both planes. Normally, only one plane is selected for readback, but where deemed necessary both planes can be read back simultaneously.

Structurally, the BPM detectors have been cut from the inner surface of the beam pipe itself. They consist of four strips of steel, each connected to an output at the upstream end. The strips are arranged more or less diagonally around the beam pipe, as in this longitudinal view:

[Future Picture]

The principle used here is basically that of a stripline detector; several variations of this type of device will soon be encountered. As our tiny and invisible but charged particles whiz by the stripline, an image charge is created in the metal. “Image charge” is just a way of saying that as a proton passes next to the strip, electrons attracted to the positive charge will be drawn to the surface of the stripline, forming an “image” of the proton; likewise, an antiproton creates a positive “image” as the electrons are repulsed. Since the particle bunches are traveling in a direction, the image charge becomes an image current traveling in the opposite direction of the beam. It the image current which is sensed at each of the four outputs:

[Future Picture]

What makes the BPMs work as position detectors is that the strength of the signal is proportional to the proximity of the beam to the stripline. The beam position is calculated by comparing the relative strengths of the four signals. (Those who understood how BPMs worked in the Main Ring are at a slight disadvantage now, because there were two dedicated plates for measuring the horizontal position, and two for the vertical. In the Main Injector, information for the two planes must be extracted from all four detectors.) Suppose that the beam is left of center. The two striplines on the left will each see a stronger signal than the two on the right.

BPM Electronics

The signal developed from the image current is first sent to an RF module, located upstairs in the service building. There is a dedicated RF module for each of the BPMs. The term "RF" may be a bit misleading in this case, because the modules have nothing to do with generating the RF for the acceleration systems. They can be better thought of as RF detectors or receivers, tuned to 53 MHz because that is the rate at which bunches are passing through the detector. They compare the signal strengths from the BPMs and convert them into DC analog signals palatable to the next stage in the processing, the Analog Box.

The Analog Box—another slightly misleading name—is a dedicated MADC that digitizes the analog signals from the RF modules and delivers the data to the microprocessor.

BPM Controls

In the Main Injector, the BPM microprocessors are actually called MBP (Main Injector Beam Position) microprocessors, because the Tevatron had already copyrighted the name "BPM". The MBP is capable of organizing the BPM data coming from the Analog Box in several different ways. Data is written into a variety of buffers. Keep in mind that the microprocessor at a given service building only handles data for the BPMs under its control, and only sees part of the picture. When a BPM application program is called up in the control room, the Main Injector front end gathers data from each of the service buildings and integrates it into a single picture, or frame. A frame represents data taken at a given moment in time.

The most fundamental variation in the way data is taken and stored is between flash frames and snapshot frames. Flash frames only look at one turn of beam, that is, a single sample of beam all the way around the ring once, or once through a beam line. They are usually used to look at either the first or last turn of beam in the machine, at the moment of injection or extraction. Flash frames require exquisite timing, which is orchestrated through the beam sync clocks. The sample of beam must be measured at each BPM location as it moves around the ring at nearly the speed of light. There is a beam sync event for the injection or extraction kickers; a delay set by a CAMAC 279 card waits for a specified number of Main Injector revolutions before triggering the BPMs. When the delay is complete, the flash trigger is fanned out simultaneously to all of the service buildings. Local delays are implemented at the service buildings to ensure that the trigger arrives precisely when the sample of beam is passing through the BPM. The flash data is then stored into a dedicated flash buffer. The data in the buffer is replaced with new data each time a new injection or extraction event is issued.

The snapshot frames do not need to be timed quite as stringently. The BPM signals are averaged over several turns—8 turns is typical—before being assembled into a frame. After this flurry of activity, there is a wait of several milliseconds before data for the next frame is taken:

[Future Picture]

The circular buffer for the snapshot frames (that is, a buffer that continuously writes new data over the oldest data) holds up to 512 entries. It “freezes” at the end of a beam cycle, or when an abort occurs, saving the last second or so of data. In that way, the history of a problem can be analyzed as it developed.

There are two other buffers that take their data from the circular buffer; these are snapshot frames that have been earmarked for specific times. One is the display buffer, for which one time in the machine cycle can be selected and displayed for every pulse. New data is written into the display frame once a cycle.

The other is the profile buffer, which selects a series of sample times within a cycle. It holds up to 128 entries, all taken from the snapshot buffer, and are spaced at equal intervals chosen by the user. It also is refreshed once a cycle.

Several additional application programs use the BPMs. They will be covered in the prophesied chapter on beam tuning.

In the beam lines the MBP microprocessors are not used. Any given pulse of beam only passes through once. Data is sent back to the appropriate front ends via the MADC chassis, and the applications software assembles the flash frame. There is no need for the relatively complex overhead required to generate the display, snapshot, and profile frames.

Beam Loss Monitors

When all else has failed, and beam is so far off course that none of the other beam detectors will work, there are always the beam loss monitors. (The BLMs are useful in less drastic situations as well.) The BLMs are placed outside the beam pipe in order to detect the spray of secondary particles created when beam strikes an object, such as the inside of a beam pipe.

Externally, a BLM presents itself as a metal cylinder, usually a few inches long, with cables connected at either end. Inside the cylinder is a sealed glass tube filled with argon. The tube also houses electrodes: the anode runs along the axis of the tube, and the cathode is a metal cylinder just inside the surface of the glass. There is a potential of about 2 KV across the electrodes. When a particle passes through the glass tube, some of the argon is ionized. The ions drift toward their favorite electrode and a current, proportional to the number of particles passing through, is produced.

The red cables connected to the BLMs provide the high voltage to the electrodes; the power supply for these is upstairs. The green cable connected on the other side of the cylinder carries the detected signal to the electronics, which is also upstairs.

There are two integrators in the electronics; the signal from the BLM is split and sent to both integrators. One of the integrators is “lossy” and provides something resembling a real-time signal; the other actually integrates.

The microprocessor that coordinates BLM data is actually the BPM microprocessor. One consequence of this union is that the times for sampling the losses correspond to the BPM timers: e.g. the BLM Display timer will be the same as the BPM Display timer.

As with the BPMs, BLMs in the beam lines return data to the controls system through the MADCs instead of MBP microprocessors.